

AUTOMATIC CLUSTERING OF TOKENS FROM A CORPUS FOR GRAMMAR ACQUISITION

PRIORITY APPLICATION

The present application claims priority to U.S. Patent Application 09/912,461, filed July 26, 2001, the contents of which are incorporated herein by reference.

BACKGROUND

The present invention relates to an application that builds linguistic models from a corpus of speech.

For a machine to comprehend speech, not only must the machine identify spoken (or typed) words, but it also must understand language grammar to comprehend the meaning of commands. Accordingly, much research has been devoted to the construction of language models that a machine may use to ascribe meaning to spoken commands. Often, language models are preprogrammed. However, such predefined models increase the costs of a speech recognition system. Also, the language models obtained therefrom have narrow applications. Unless a programmer predefines the language model to

recognize a certain command, the speech recognition system that uses the model may not recognize the command. What is needed is a training system that automatically extracts grammatical relationships from a predefined corpus of speech.

SUMMARY

An embodiment of the present invention provides a method of learning grammar from a corpus, in which context words are identified from a corpus. For the other non-context words, the method counts the occurrence of predetermined relationships with the context words, and maps the counted occurrences to a multidimensional frequency space. Clusters are grown from the frequency vectors. The clusters represent classes of words; words in the same cluster possess the same lexical significance and provide an indicator of grammatical structure.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a flow diagram of a method of an embodiment of the present invention.

FIG. 2 illustrates mapping frequency vectors that may be obtained during operation of the present invention.

FIG. 3 illustrates an exemplary cluster tree.

DETAILED DESCRIPTION

Embodiments of the present invention provide a system that automatically builds a grammatical model from a corpus of speech. The present invention uses clustering to group words and/or phrases according to their lexical significance. Relationships between high frequency words called Acontext words@ and other input words are identified. The words to be clustered are each represented as a feature vector constructed from the identified relationships. Similarities between two input words are measured in terms of the distance between their feature vectors. Using these distances, input words are clustered

according to a hierarchy. The hierarchy is then cut at a certain depth to produce clusters which are then ranked by a "goodness" metric. Those clusters that remain identify words or tokens from the corpus that possess similar grammatical significance.

Clustering per se is known. In the context of language modeling, clustering has typically been used on words to induce classes that are then used to predict smoothed probabilities of occurrence for rare or unseen events in the training corpus. Most clustering schemes use the average entropy reduction to decide when two words fall into the same cluster. Prior use of clustering, however, does not provide insight into language model of grammar.

FIG. 1 illustrates a method of the present invention according to a first embodiment. The method operates upon input text, a set of words from which the grammar model shall be constructed. Typically, the input text comprises a set of single words or phonemes. From the input text, the method identifies context words (Step 1010). Context words are those words or phonemes in the input text that occur with the highest frequency. The method 1000 may cause a predetermined number of words (say, 50) that occur with the highest frequency to be identified as context words.

The method 1000 determines relationships that may exist between the context words and the remaining words, called "input words" herein, in the input text. For example, the method 1000 may determine how many times and in which positions an input word appears adjacent to a context word. Table 1 below illustrates relationships that may exist between certain exemplary input words and exemplary context words.

	Context Word											
	to				from				in			
Input Word	-2	-1	1	2	-2	-1	1	2	-2	-1	1	2
Chicago	f_{111}	f_{112}	f_{113}	f_{114}	f_{121}	f_{122}	f_{123}	f_{124}	f_{131}	f_{132}	f_{133}	f_{134}
New York	f_{211}	f_{212}	f_{213}	f_{214}	f_{221}	f_{222}	f_{223}	f_{224}	f_{231}	f_{232}	f_{233}	f_{234}
Baltimore	f_{311}	f_{312}	f_{313}	f_{314}	f_{321}	f_{322}	f_{323}	f_{324}	f_{331}	f_{332}	f_{333}	f_{334}
red	f_{411}	f_{412}	f_{413}	f_{414}	f_{421}	f_{422}	f_{423}	f_{424}	f_{431}	f_{432}	f_{433}	f_{434}
white	f_{511}	f_{512}	f_{513}	f_{514}	f_{521}	f_{522}	f_{523}	f_{524}	f_{531}	f_{532}	f_{533}	f_{534}
blue	f_{611}	f_{612}	f_{613}	f_{614}	f_{621}	f_{622}	f_{623}	f_{624}	f_{631}	f_{632}	f_{633}	f_{634}

Table 1

Each entry of the table, f_{ijk} represents, for a given input word i ; how many times a context word C ; and non-context word j ; appears within a predetermined relationship. Thus, f_{111} - f_{114} each represent the number of times the input word "Chicago" and the context word "to" appear within adjacencies of -2 words, -1 word, +1 word and +2 words respectively.

Based upon the frequencies, an N dimensional vector may be built for each input word (step 1020). The number of dimensions N of the frequency vector is a multiple of the total number of context words, the total number of input words and the total number of relations identified by the method 1000. The vector represents grammatical links that exist between the input words and the context words. Thus, each input word maps to an N dimensional frequency space. A representative frequency space is shown in FIG. 2 ($N=3$).

The method 1000 builds clusters of input words (Step 1030). According to the principles of the present invention, input words having the same lexical significance should possess similar vectors in the frequency space. Thus, it is expected that city names will

exhibit frequency characteristics that are similar to each other but different from other input words having a different lexical significance. They will be included in a cluster (say, cluster 10, FIG. 2). So, too, with colors. They will be included in another cluster (say, cluster 20). Where words exhibit similar frequency significance, they are included within a single cluster.

As is known, a cluster may be represented in an N-dimensional frequency space by a centroid coordinate and a radius indicating the volume of the cluster. The radius indicates the "compactness" of the elements within a cluster. Where a cluster has a small radius, it indicates that the elements therein exhibit a very close relationship to each other in the frequency space. A larger radius indicates fewer similarities between elements in the frequency space.

The similarity between two words may be measured using the Manhattan distance metric between their feature vectors. Manhattan distance is based on the sum of the absolute value of the differences among the vector's coordinates. Alternatively, Euclidean and maximum metrics may be used to measure distances. Experimentally, the Manhattan distance metric was shown to provide better results than the Euclidean or maximum distance metrics.

Step 1030 may be applied recursively to grow clusters from clusters. That is, when two clusters are located close to one another in the N dimensional space, the method 1000 may enclose them in a single cluster having its own centroid and radius. The method 1000 determines a distance between two clusters by determining the distance between their centroids using one of the metrics discussed above with respect to the vectors of input words. Thus, the Manhattan, Euclidean and maximum distance metrics may be used.

A hierarchical "cluster tree" is grown representing a hierarchy of the clusters. At one node in the tree, the centroid and radius of a first cluster is stored. Two branches extend from the node to other nodes where the centroids and radii of subsumed clusters are

stored. Thus, the tree structure maintains the centroid and radius of every cluster built according to Step 1030. Step 1030 recurs until a single, all encompassing cluster encloses all clusters and input words. This cluster is termed the "root cluster" because it is stored as the root node of the cluster tree. An exemplary cluster tree is shown in FIG. 3.

As will be appreciated, the root cluster N13 has a radius large enough to enclose all clusters and input words. The root cluster, therefore, possesses very little lexical significance. By contrast, "leaf clusters," those provided at the ends of branches in the cluster tree, possess very strong lexical significance.

At Step 1040, the method 1000 cuts the cluster tree along a predetermined line in the tree structure. The cutting line separates large clusters from smaller clusters. The large clusters are discarded. What remains are smaller clusters, those with greater lexical significance.

The cutting line determines the number of clusters that will remain. One may use the median of the distances between clusters merged at the successive stages as a basis for the cutting line and prune the cluster tree at the point where cluster distances exceed this median value. Clusters are defined by the structure of the tree above the cutoff point.

Finally, the method 1000 ranks the remaining clusters (Step 1050). The lexical significance of a particular cluster is measured by its compactness value. The compactness value of a cluster simply may be its radius or an average distance of the members of the cluster from the centroid of the cluster. Thus, the tighter clusters exhibiting greater lexical significance will occur first in the ranked list of clusters and those exhibiting lesser lexical significance will occur later in the list. The list of clusters obtained from Step 1050 is a grammatical model of the input text.

The method 1000 is general in that it can be used to cluster "tokens" at any lexical level. For example, it may be applied to words and/or phrases. Table 2 illustrates the result of clustering words and Table 3 illustrates the result of clustering phrases as performed on

an experimental set of training data taken from the How May I Help You? Training corpus disclosed in Gorin, et al., "How May I Help You?," vol. 23, Speech Communication, pp. 113-127 (1997). Other lexical granularities (syllables, phonemes) also may be used.

Class Index	Compactness Value	Class Members
C363	0.131	make place
		eight eighty five four nine oh one seven six three two
C118	0.18	zero
C357	0.19	bill charge
C260	0.216	an and because but so when
C300	0.233	K O ok
C301	0.236	From please
C277	0.241	again here
C202	0.252	as it's
C204	0.263	different third
C77	0.268	number numbers
C275	0.272	Need needed want wanted
C256	0.274	assistance directory information
C197	0.278	all before happened
C68	0.278	ninety sixty
C41	0.29	his our the their
C199	0.291	called dialed got have
		as by in no not now of or something that that's there
C27	0.296	whatever working
C327	0.296	I I'm I've
		canada england france germany israel italy japan
C48	0.299	london mexico paris
C69	0.308	back direct out through
C143	0.312	connected going it
		arizona california carolina florida georgia illinois island
		jersey maryland michigan missouri ohio pennsylvania
C89	0.314	virginia west york
C23	0.323	be either go see somebody them
C90	0.332	about me off some up you

Table 2: Results of Clustering Words from AT&T's *How May I Help You ?* Corpus

Class Index	Compactness Value	Class Members
D365	0.226	wrong:C77 second
D325	0.232	C256:C256 C256
D380	0.239	area:code:C118:C118:C118:C118:C118 C68
D386	0.243	A:C77 this:C77
D382	0.276	C260:C357:C143:to:another C260:C357:C143:to:my:home
D288	0.281	C327:C275:to:C363 I'd:like:toC363 to:363 yes:I'd:like:to:C363
D186	0.288	good:morning yes:ma'am yes:operator hello hi ma'am may well
D148	0.315	problems trouble
D87	0.315	A:T:C260:T C260:C327 C27:C27 C41:C77 C118 C143 C260 C197 C199 C202 C23 C260 C27 C277 C301 C69 C77 C90 operator to
D183	0.321	C118:C118:hundred C204 telephone
D143	0.326	new:C89 C48 C89 colorado massachusetts tennessee texas
D387	0.327	my:home my:home:phone
D4	0.336	my:calling my:calling:card my:card
D70	0.338	C199:a:wrong:C77 misdialed
D383	0.341	like:to:C363 trying:to:C363 would:like:to:C363
D381	0.347	like:to:C363:a:collect:call:to like:to:C363:collect:call would:like:to:C363:a:collect:call would:like:to:C363:a:collect:call would:like:to:C363:a:collect:call:to
D159	0.347	C118:C118 C118:C118:C118 C118:C118:C118:C118:C118:C118

Class Index	Compactness Value	Class Members
		C118:C118:C118:C118:C118:C118:C118
		C118:C118:C118:C118:C118:C118:C118:C118:C118:C118
		C:118:C118:C118:C118:C118:C118:C118:C118:C118:C118:C118
		8 area:code:C118:C118:C118 C300

Table 3: Results from a First Iteration of Combining Phrase Acquisition and Clustering from the *How May I Help You?* Corpus (Words in a Phrase are Separated by a Colon).

Adjacency of words is but one relationship that the method 1000 may be applied to recognize from a corpus. More generally, however, the method 1000 may be used to recognize predetermined relationships among tokens of the corpus. For example, the method 1000 can be configured to recognize words that appear together in the same sentences or words that appear within predetermined positional relationships with punctuation. Taken even further, the method 1000 may be configured to recognize predetermined grammatical constructs of language, such as subjects and/or objects of verbs. Each of these latter examples of relationships may require that the method be pre-configured to recognize the grammatical constructs.

Several embodiments of the present invention are specifically illustrated and described herein. However, it will be appreciated that modifications and variations of the present invention are covered by the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention.